ADAPTIVE ROUTING ON THE PLANE

LEANDROS TASSIULAS

University of Maryland, College Park, Maryland

(Received December 1992; revisions received August 1994; accepted January 1995)

Demands for service arrive at random times, in random locations, in a region of the plane. The service time of each demand is random. A server that travels with constant speed moves from demand to demand providing service. The server spends its time either in providing service or in traveling. The objective is to route the server, based on the location of the current demands on the plane and the anticipated demand arrivals, such that the time spent in traveling is minimal and the service is provided efficiently. A routing policy is provided that achieves maximum throughput and is independent of the statistical parameters of the system, under the assumption that the arrival process is Poisson. For a renewal arrival process, a class of policies is specified that achieve maximum throughput based on some knowledge of the system parameters. Finally, an adaptive version of the partitioning policies is given, which makes them independent of the system statistics.

The problem of routing in the Euclidian plane has a long history. The most heavily studied problem in this context is the Euclidean Traveling Salesman Problem (ETSP), the version of the Traveling Salesman Problem where the objective is to find the route that passes through every point of a group of points located in the plane, and has minimum length. Both the great practical importance of the ETSP as well as its mathematical difficulty justify the vast attention that it attracted in the past; some of the previous work is reported in Karp (1977). Different optimization objectives than the total traveling time have also been considered in the past. One of them is the average waiting time for each point, where the waiting time of a point is the elapsed time from the beginning until the time that the server visits the point and the average is over all points. The latter objective is more appropriate in situations where the points correspond to the sites of service demands which are satisfied by a server moving from demand to demand.

The basic characteristics of the above problems are that they are deterministic and static. The number of points to be visited as well as their locations are predetermined and fixed. Also the service times are known. For several practical problems of this nature the number of points requiring service is not fixed since new demands arrive during the time that the server is providing service. Furthermore the locations as well as the service times are not known in advance and have to be modeled as random quantities. The nature of these problems can be captured closer by a dynamic stochastic model. Recently there has been considerable research activity in this direction. Models that capture the random and dynamic nature of the problem were considered in Psaraftis (1988), Jaillet (1988), Batta et al. (1988), Bertsimas and van Ryzin (1991, 1993a, b), Thompson and Psaraftis (1993). Psaraftis considered a model where the requests are located at the nodes of a graph and the server can move from any node to any other node, spending a traveling time depended on the nodes; the problem of minimizing the average waiting time was addressed. Batta et al. considered a model where requests for service arise at the nodes of a graph and they are served by a server which after serving each request needs to return to its *home* position. The problem is then selecting that home position. Jaillet studied a version of the Traveling Salesman Problem where a random subset of the nodes have to visited. More recently, Bertsimas and van Ryzin (1991) considered a model where the service demands arise at any point of a bounded convex region of the plane. The locations of the requests constitute an i.i.d. process with uniform distribution on the plane. The server can move from request to request in a time proportional to the physical distance of the points. The delay and the throughput of several routing policies were analyzed. The multivehicle case and the case of general arrival and location processes have also been considered by Bertsimas and van Ryzin (1993a, b).

In this paper a model similar to the one in Bertsimas and van Ryzin (1991) is considered. The service requests may arise at any point in a region of the plane. The distribution of the requests can be arbitrary, and the locations of different requests can be correlated. The objective is to obtain routing policies which achieve large throughput, do not rely on information about the statistics of the system, and meet the performance objectives for a wide range of arrivals and service rates.

A routing policy determines the order in which the server visits the service requests; the routing decisions may depend on the configuration of the requests on the plane, possibly in a complicated manner. The routing policies presented here are based on an algorithm for computing a route that goes through a number of points on the plane and keeps the total traveling distance short. The algorithm, called *congestion focusing algorithm* (CFA), computes the route in a recursive manner; in each step the server focuses in a region of the service area and selects the next

Subject classifications: Queues: algorithms: Markovian. Network/graphs: Routing, stochastic traveling salesman. Transportation: vehicle routing.
Area of review: DISTRIBUTION, TRANSPORTATION, AND LOGISTICS.

0030-364X/96/4405-0823 $01.25

demand to be served from this region. After the demand is served the server focuses in the subregion with the highest concentration of demands and selects the next demand to be served from there. The demand is selected such that the server moves toward the area where the concentration of the service demands is increasing. The policy routes the server with repetitive application of the CFA algorithm.

If the system is stable, then the throughput is equal to the arrival rate. The maximum arrival rate for which the system is stable under a particular policy π is equal to the maximum achievable throughput under π, and it will be referred as the throughput of π in the following. A necessary stability condition for every policy is that the arrival rate is strictly smaller than the inverse of the expected service time. The policy based on *CFA* has the property of achieving maximum throughput in the sense that it stabilizes the system for all arrival rates strictly smaller than the inverse of the expected service time. In addition, this policy has two advantages over other maximum throughput policies that have been proposed previously: it does not need knowledge of the arrival rates for its implementation; and it preserves the maximum throughput properties for all sample paths of demand locations, that is independently of the distribution and the dependencies of the locations on the plane. If the arrival rate and the distribution of the locations is known, then the comparison of the performance of different routing policies is on the basis of the expected waiting time they achieve. If the arrival rate and the distribution of the locations is not specified, then we cannot always compare two policies since it is possible that none of them achieves smaller waiting time than the other, uniformly over all distributions of locations and arrival rates. Usually in applications, either the arrival rate or the distribution of the locations, or both are unknown. In this case a policy that guarantees the stability, like the one proposed here, is preferable over policies which are either parametrized by the arrival rate and/or may be unstable for certain distributions of locations.

The routing policy based on the CFA algorithm is analyzed for Poisson arrivals, and the properties mentioned above are shown. A parametrized class of policies which stabilize the system for general arrivals is specified next. In order to stabilize the system the parameter is selected based on the arrival and service rates. Finally an adaptive version of the partitioning policies considered by Bertsimas and van Ryzin (1991) is given. The adaptive policy selects the resolution of the partition adaptively, and hence it becomes independent of the statistics of the system. Unlike the previous work on the subject, in our study the throughput results are obtained for every sample path of the locations of the arriving demands. Therefore no assumption on the distribution of the demands on the plane is needed, and the location of different demands can be arbitrarily dependent.

The rest of the paper is organized as follows. In Section 1 the model and the performance objectives are defined and some notation is introduced. In Section 2 the CFA is given. In Section 3 the maximum throughput policy is given for Poisson arrivals. In Section 4 the case of general arrivals is considered and the class of stabilizing policies is specified. In Section 5 the adaptive partitioning policies are considered. Finally, in Section 6 the results are summarized, and open problems for further research are identified.

1. THE MODEL

We consider a region $\mathcal{A}$, which is a bounded convex subset of R^2 and arbitrary otherwise. Demands for service are generated at random times and in random locations in $\mathcal{A}$. The demands are generated according to an arrival process $A = \{A(t), t \geq 0\}$ where $A(t)$ is the number of demands that have been generated up to time t. Some of the results in the paper are obtained under the assumption that A is Poisson and some others under the more general assumption that A is a renewal process; the specific statistical assumptions on A are mentioned as needed later. Under both assumptions we denote by λ^{-1} the expected interarrival time. The service time of a demand is random. Let S_i be the service time of the ith demand. We assume that $S = \{S_i\}_{i=1}^{\infty}$ is an i.i.d. process and denote the expected service time $E[S_i]$ by μ^{-1}. Let (X_i, Y_i) be the coordinates of the ith demand. We make no assumption about the statistics of the *location process* $\{(X_i, Y_i)\}_{i=1}^{\infty}$; in fact the results we obtain hold for every sample path of the location process.

The server moves on the plane with unit speed: hence in the rest of the paper the distance and the traveling time between two points are used interchangably. We assume that the server moves from demand to demand in the straight line segment connecting the locations of the two demands. Hence the route of the server is a polygonal line and is specified completely by the sequence of locations that the server visits. A polygonal line that connects a number of points is called *tour* in the following and is represented by the sequence of the points in the order that they appear in the line.

A routing policy is any rule for selecting the order in which the server will serve the demands. Let $Q(t)$ be the number of demands on the plane at time t including the demand under service at that time, if any. The system is *stable* if

$$\limsup_{t\to\infty} \frac{1}{t} \int_0^t EQ(s)\, ds < \infty .$$

Clearly if the system is stable according to the above definition the long run average delay experienced by the demands is finite from Little's law, and the throughput of the system is equal to the arrival rate. In the rest of the paper we consider the normalized system throughput that is $\rho = \lambda/\mu$. The throughput of a policy is defined as the supremum of ρ over all the arrival rates for which the system is stable. Clearly for any $\rho \geq 1$ the system is unstable, hence the throughput of a policy is always less than or equal to one. The routing algorithm of the following section is employed to design policies that indeed achieve throughput 1, with the understanding that every throughput strictly less than one is indeed achievable, while for $\rho = 1$ the system is unstable and the throughput $\rho = 1$ is not achievable.

2. THE CONGESTION FOCUSING ALGORITHM

The recursive algorithm presented below produces a tour in the plane that has the following property. For any arrangement of N points on a convex bounded region of the plane, there is a route that goes through at least $\log N/2$ of these points and the total traveling distance is bounded above by a number independent of N (log denotes the logarithm with base 2; if $\log N$ is an odd number, then $\log N/2$ stands for the largest integer which is smaller than $\log N/2$). For the description of the algorithm we consider that region to be enclosed within a square with edge of length a. This is always possible if we take a to be the maximum distance between any two points in $\mathcal{A}$. The algorithm acts as follows. It starts from a point in region $\mathcal{A}$: then it separates the square that surrounds the region in four equal squares (Figure 1), and it selects the one with the largest number of points. It moves to a point in that square and the same process is repeated: that is the latter square is divided in four equal squares and the one with the largest number of points is selected and so on. The algorithm is specified next.

Congestion Focusing Algorithm

Input. A square A_0 with edge length a enclosing a convex region of the Euclidean plane with N points on it and a starting point p_0.

Output. A polygonal line specified by the sequence of distinct points $p_0, p_1, p_2, \ldots, p_K$

STEP 0. Set $i = 0$, exclude point p_0 from the configuration.

STEP 1. Divide the square A_i in four equal squares as indicated in Figure 1. Let A_{i+1} be the one of these small squares with the largest number of points (ties are broken arbitrarily).

STEP 2. If A_{i+1} is empty then stop. Otherwise select an arbitrary point of this square, to be point p_{i+1} and exclude it from the configuration.

STEP 3. Set $i = i + 1$ and go to step 1.

A route selected by the algorithm is drawn in Figure 2.

Theorem 1. *Consider a configuration of N points located arbitrarily on a convex bounded region of the plane which is enclosed in a square with edge length equal to a. The number K of points visited by CFA upon termination, when we start from an arbitrary initial point, is greater than or equal to $\log N/2$. Furthermore the length of the polygonal line defined by these points is less than or equal to $2\sqrt{2}a$.*

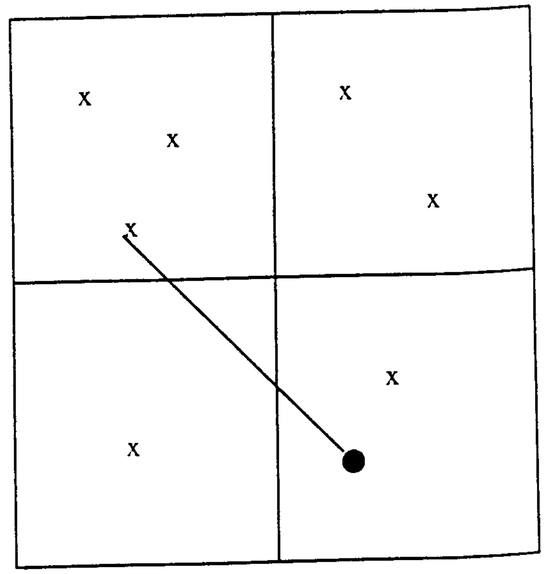

Figure 1. In each iteration of Algorithm A1 the square A_i is divided in four pieces as in the above picture. A point is selected for service in the square with the maximum number of points. The later square is divided in four pieces again, and the process is repeated.

Proof. We show first that if $i \leq \log N/2$ then the number of points in the square A_i in the ith step of CFA is greater than or equal to $(N - \sum_{l=0}^{i-1} 4^l)/4^i > 0$. We show that by induction. For $i = 0$ this is true from the initial condition of the algorithm. Assume that it is true for some $i > 0$ and less than $\log N/2$. We will show that it is true for $i + 1$ as well. From the induction hypothesis we have that after the deletion of point p_i the number of points in the square A_i is equal to $(N - \sum_{l=0}^{i-1} 4^l)/4^i - 1 = (N - \sum_{l=0}^{i} 4^l)/4^i$. The square A_{i+1} is obtained by dividing A_i in four pieces and taking the one with the largest number of points; therefore

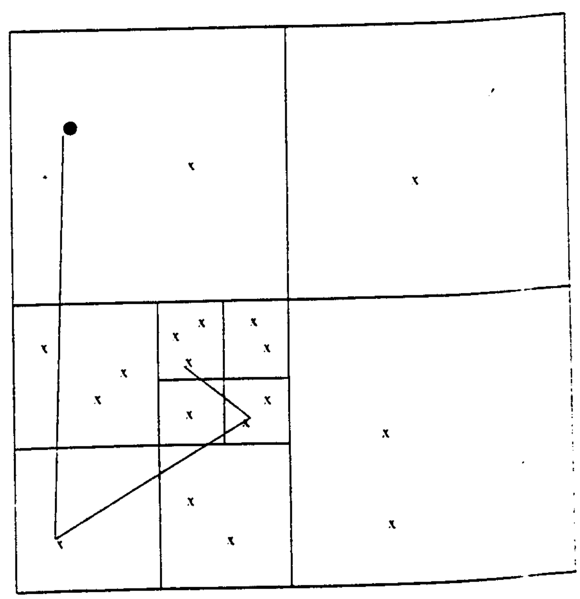

Figure 2. The tour computed by Algorithm A1 for the configuration of points in the picture is depicted.

it contains at least one-fourth of the points in A_i which is $(N - \Sigma_{l=0}^{i} 4^l)/4^{i+1}$. Hence as long as $i < \log N/2$ the square A_i is nonempty, the point p_i is well defined, and the algorithm proceeds to the next step. Therefore the number of points selected by the algorithm is greater than or equal to $\log N/2$ as it is stated in the theorem.

We show now that the length d of the tour connecting the points $p_0, \ldots, p_K$ produced by CFA is less than or equal to $2\sqrt{2}a$. The points p_i, p_{i-1} are contained in the ith square considered by CFA which has edge length equal to $a/4^i$; therefore the line than connects them will have length $d_{p_i p_{i-1}}$ that is less than or equal to the length of the diagonal of this square, which is $\sqrt{2}a/4^i$. Therefore the total length of the tour will be

$$d = \sum_{i=0}^{K} d_{p_i p_{i-1}} \leq \sum_{i=0}^{K} \sqrt{2}\frac{a}{2^i} \leq \sum_{i=0}^{\infty} \sqrt{2}\frac{a}{2^i} = 2\sqrt{2}a . \quad \square$$

The intuition behind the algorithm is the following. It is clear that if the number of points which are uniformly allocated in a fixed region of the plane increases then, roughly speaking, the density of points in every location of the region increases as well, and the distance between them decreases. If the points are not arranged uniformly but in an arbitrary manner then still we expect that as the number of points increases, in certain locations the points will come closer one to the other. Hence if the server moves in the areas with high concentration of points the time wasted in traveling is reduced. The algorithm presented in this section captures that intuition: the server continuously approaches the location with the high concentration of points in a recursive manner. It does so indeed since as we show the number of points visited by the algorithm increases as the number of points in $\mathcal{A}$ increases while the length of the tour remains bounded.

3. THE CONGESTION FOCUSING POLICY

We specify next a routing policy, called *congestion focusing policy* (CFP), that is based on CFA. That routing policy stabilizes the system for all $\rho < 1$, i.e., it achieves maximum throughput. It does so under the following statistical assumption.

S1. The arrival process A is Poisson, the service times are i.i.d., and they have finite second moments.

The policy CFP acts as follows. At certain time instants τ_i, $i = 1, \ldots,$ CFP selects a sequence of demands to be served in the time interval $[\tau_i, \tau_{i+1})$ as well as the route that the server should follow. The demands are among those which are on the plane at time τ_i. The decision is taken based on the configuration of the points in region $\mathcal{A}$ at that time instant.

Congestion Focusing Policy

Initialize. $\tau_0 = 0$, $i = 0$.

STEP 1. At time τ_i a sequence of demands is determined by CFA for the configuration of the demands in $\mathcal{A}$ at that time. From theorem 3.1 these demands are more than or equal to $\log Q(\tau_i)/2$. Select for service the first $\log Q(\tau_i)/2$ of these demands.

STEP 2. The demands selected in Step 1 are served in the order that they have been selected. If after the end of service the total time that has been spent in traveling since time τ_i is less than $2\sqrt{2}a$, the server idles for an amount of time such that the total time spent in idling and traveling is equal to $2\sqrt{2}a$. The time at which the idling period ends is the next decision time instant τ_{i+1} at which the next sequence of points is selected. Set $i = i + 1$ and go to step 1.

Remarks

1. Note that if at the end of an idling period the plane is empty, then the new cycle will include no service but only idling of duration $2\sqrt{2}a$.
2. The idling in Step 2 of CFP is introduced for technical reasons only such that the process of the number of demands is a Markov chain. The Markovian property of the demand process will be essential for showing the stability of the policy. It seems, though, that the idling has no real significance for the operation of the policy and can be omitted in an implementation of the policy.

In the rest of this section we show that CFP achieves maximum throughput as it is stated in the following theorem.

Theorem 2. *When the system is operated under CFP and the statistics satisfy assumption S1 then for every $\rho < 1$*

$$\limsup_{t\to\infty} \frac{1}{t}\int_0^t EQ(s)\,ds < \infty , \tag{1}$$

and the system is stable.

In order to prove Theorem 2 we study first the behavior of the process $\mathbf{Q} = \{Q_i\}_{i=1}^{\infty}$, where $Q_i = Q(\tau_i)$. In the next lemma we show a property of the drift of the square of the process $\mathbf{Q}$ that is used in the proof of Theorem 2.

Lemma 1. *When the system is operated under CFP, $\rho < 1$ and the statistics satisfy assumption A1 then there exist B, $\epsilon > 0$ such that*

$$E[V(Q_{i+1}) - V(Q_i)|Q_i] < -\epsilon Q_i \log Q_i \quad \text{if } Q_i > B , \tag{2}$$

where $V(Q) = Q^2$.

Proof. After some algebra we have

$$E[V(Q_{i+1}) - V(Q_i)|Q_i]$$
$$= E[(Q_{i+1} - Q_i)^2|Q_i] + 2Q_i E[Q_{i+1} - Q_i|Q_i] . \tag{3}$$

From CFP we have

$$Q_{i+1} - Q_i = A(\tau_{i+1}) - A(\tau_i) - \frac{\log Q_i}{2} , \tag{4}$$

from which we get for the first term in the right side of (3) that

$$E[(Q_{i+1} - Q_i)^2|Q_i] = E[(A(\tau_{i-1}) - A(\tau_i))^2|Q_i]$$
$$- \log Q_i E[A(\tau_{i+1}) - A(\tau_i)|Q_i] + \frac{\log^2 Q_i}{4}$$
$$\leq \lambda^2 E[(\tau_{i+1} - \tau_i)^2|Q_i] + \frac{\log^2 Q_i}{4}, \quad (5)$$

since the arrivals are Poisson and $A(\tau_{i+1}) - A(\tau_i)$ is nonnegative. Notice that $\tau_{i+1} - \tau_i$ is equal to the total time spent in the service of the $\log Q_i/2$ demands plus the time spent in traveling from demand to demand and in idling. Since the points are selected using A1 and from Theorem 2 we get that the total time spent in traveling plus idling is equal to $2\sqrt{2}a$. Let S_j be the service time of the jth customer served in the time interval $\tau_{i+1} - \tau_i$. Using the inequality $(\Sigma_{i=1}^N a_i)^2 \leq N(\Sigma_{i=1}^N a_i^2)$ we get

$$E[(\tau_{i+1} - \tau_i)^2|Q_i]$$
$$\leq E\left[\left(\frac{\log Q_i}{2} + 1\right)\left(\sum_{j=1}^{\log Q_i/2} S_j^2 + 8a^2\right)\Big| Q_i\right]$$
$$= \left(\frac{\log Q_i}{2} + 1\right)\left(\frac{\log Q_i}{2} E[S_j^2] + 8a^2\right). \quad (6)$$

Also we have

$$E[\tau_{i+1} - \tau_i|Q_i] = \mu^{-1}\frac{\log Q_i}{2} + 2\sqrt{2}a, \quad (7)$$

since in Step 2 of the algorithm we let the server to spent exactly $2\sqrt{2}a$ time in travel plus idling. For the second term in the right side of (3) we get using (7)

$$2Q_i E[Q_{i+1} - Q_i|Q_i] = 2Q_i E[A(\tau_{i+1}) - A(\tau_i)]$$
$$- 2Q_i \frac{\log Q_i}{2}$$
$$= 2Q_i \lambda E[\tau_{i+1} - \tau_i]$$
$$- 2Q_i \frac{\log Q_i}{2} \quad (8)$$
$$\leq (\rho - 1)Q_i \log Q_i + 4\sqrt{2}aQ_i.$$

By replacing in (3) from (5), (6), and (8) we get

$$E[V(Q_{i+1}) - V(Q_i)|Q_i] \leq (\rho - 1)Q_i \log Q_i + c_1 Q_i$$
$$+ c_2 \log Q_i + c_3 \log^2 Q_i. \quad (9)$$

for certain constants c_1, c_2, c_3. In the right side of (9) the term $Q_i \log Q_i$ clearly dominates for large Q_i; hence it is clear that for every ϵ, $0 < \epsilon < (1 - \rho)$ there exists B such that (2) holds. □

We can easily see that the chain **Q** is irreducible and aperiodic, therefore using Foster's criterion, Asmussen (1987), we get the following.

Corollary 1. *The Markov chain* **Q** *is positive recurrent.*

Consider the times T_i, $i = 0, 1, \ldots$ at which the system empties for ith time. That is

$$T_0 = 0,$$
$$T_i = \min\{\tau_j: \tau_j > T_{i-1}, Q_j = 0\}, \quad i = 1, \ldots.$$

The times T_i are well defined for all i's since from Corollary 1 the system will empty infinitely often almost surely. The following lemma is used in the proof of the theorem.

Lemma 2.

$$E\int_0^{T_1} Q(s)\, ds < \infty.$$

Proof. With $\mathcal{G}_n$ denoting the sigma field generated by $(Q_k, k \leq n)$, $n = 1, 2, \ldots$, let

$$\nu := \inf\{n \geq 1: Q_n = 0\},$$

which is a $\mathcal{G}_n$-stopping time. Observe that we can write

$$\int_0^{T_1} Q(s)ds \leq \sum_{k=0}^{\nu-1} (\tau_{k+1} - \tau_k)$$
$$\cdot [Q(\tau_k) + A(\tau_{k+1}) - A(\tau_k)]; \quad (10)$$

so that it suffices to show

$$E\sum_{k=0}^{\nu-1} (\tau_{k+1} - \tau_k)Q(\tau_k) < \infty, \quad (11)$$

and

$$E\sum_{k=0}^{\nu-1} (\tau_{k+1} - \tau_k)(A(\tau_{k+1}) - A(\tau_k)) < \infty. \quad (12)$$

By using successively the facts that $\{\nu \geq n\} \in \mathcal{G}_{n-1}$, the process $\{Q_n\}_{n=0}^{\infty}$ is Markov, and relation (7) we get

$$E\left[\sum_{k=0}^{\nu-1} (\tau_{k+1} - \tau_k)Q_k\right]$$
$$= \sum_{k=0}^{\infty} E[E[(\tau_{k+1} - \tau_k)Q_k|\mathcal{G}_k]1\{k \leq \nu - 1\}]$$
$$= \sum_{k=0}^{\infty} E[Q_k E[(\tau_{k+1} - \tau_k)|\mathcal{G}_k]1\{k \leq \nu - 1\}] \quad (13)$$
$$= \sum_{k=0}^{\infty} E[(\mu^{-1}Q_k \log Q_k + Q_k 2\sqrt{2}a)1\{k \leq \nu - 1\}]$$
$$= E\sum_{k=0}^{\nu-1} \mu^{-1}Q_k \log Q_k + E\sum_{k=0}^{\nu-1} Q_k 2\sqrt{2}a.$$

Also by using successively the facts that $\{\nu \geq n\} \in \mathcal{G}_{n-1}$, the process $\{Q_n\}_{n=0}^{\infty}$ is Markov, and relation (6) we get

$$E\sum_{k=0}^{\nu-1}(\tau_{k+1}-\tau_k)(A(\tau_{k+1})-A(\tau_k))$$
$$=\sum_{k=0}^{\infty}E[E[(\tau_{k+1}-\tau_k)(A(\tau_{k+1})-A(\tau_k))|\mathcal{G}_k]1 \cdot\{k\le\nu-1\}]$$
$$=\sum_{k=0}^{\infty}E[E[(\tau_{k+1}-\tau_k)(A(\tau_{k+1})-A(\tau_k))]|Q_k]1 \cdot\{k\le\nu-1\}]$$
$$=\sum_{k=0}^{\infty}\lambda^2E[E[(\tau_{k+1}-\tau_k)^2|Q_k]1\{k\le\nu^F-1\}]$$
$$\le E\sum_{k=0}^{\nu-1}\left(\frac{\log Q_k}{2}+1\right)\left(\frac{\log Q_k}{2}E[S_j^2]+8a^2\right). \quad (14)$$

From (13) and (14) it is clear that in order to show (11) and (12) it is enough to show

$$E\sum_{k=0}^{\nu-1}\mu^{-1}Q_k\log Q_k<\infty. \quad (15)$$

By using successively the facts $\{\nu\ge n\}\in\mathcal{G}_{n-1}$, the process $\{Q_n\}_{n=0}^{\infty}$ is Markov, Lemma 1 and $\{\nu\ge n\}\subset\{\nu\ge n-1\}$, we obtain

$$\begin{aligned}E[V(Q_n)1\{\nu\ge n\}]&=E[E[V(Q_n)|\mathcal{G}_{n-1}]1\{\nu\ge n\}]\\&=E[E[V(Q_n)|Q(\tau_{n-1})]1\{\nu\ge n\}]\\&\le E[V(Q_{n-1})1\{\nu\ge n\}]\\&\quad-\epsilon E[Q_{n-1}\log Q_{n-1}1\{\nu\ge n\}].\end{aligned} \quad (16)$$

Upon iteration, (16) yields

$$0\le E[V(Q_n)1\{\nu\ge n\}]\le EV(Q(\tau_1)) -\epsilon E\left[\sum_{k=1}^{n-1}Q_k\log Q_k1\{k\le\nu-1\}\right],$$

which implies

$$E\left[\sum_{k=1}^{\infty}Q_k\log Q_k1\{k\le\nu-1\}\right]\le\epsilon^{-1}EV(Q_1),$$

and (15) follows. □

Now we can proceed to the proof of Theorem 2.

Proof of Theorem 2. We set $\gamma(t):=\inf\{n\ge1: T_n>t\}$ and from the fact $\{\gamma(t)\ge n+1\}=\{T_n\le t\}\in\mathcal{G}_{T_n}$, we have

$$E\int_0^t Q(s)ds\le E\int_0^{T_{\gamma(t)}}Q(s)\,ds$$
$$=E\sum_{k=0}^{\infty}\int_{T_k}^{T_{k+1}}Q(s)\,ds1\{k\le\gamma(t)-1\}$$
$$=\sum_{k=0}^{\infty}E\left[E\left[\int_{T_k}^{T_{k+1}}Q(s)\,ds\middle|\mathcal{G}_{T_k}\right]1\{k\le\gamma(t)-1\}\right]$$
$$\le C_1E\gamma(t),$$

where

$$C_1:=E\int_0^{T_1}Q(s)\,ds,$$

is finite because of Lemma 2. To complete the proof, it suffices to show that for some constant c_2, $E\gamma(t)\le c_2t$, $t\ge0$. This is clearly true for $c_2=(2\sqrt{2}a)^{-1}$. □

From Theorem 2 we conclude that CFP indeed achieves maximum throughput. Also CFP does not rely on the knowledge of the arrival and service rates for the selection of the route but only on the location of the demands on the plane at the decision time instances. Hence this policy is robust on variations of the arrival and service rates during the operation of the system.

4. GENERAL ARRIVALS

Up to this point we have assumed that the process of arrivals in the system is Poisson. As a consequence of the Poisson arrivals, the process of the number of demands in the system at the decision time instants has the Markov property, which has played an essential role in obtaining the stability result in the previous section. In this section we obtain some stability results for the system when the arrival process is renewal. We present a family of routing policies $\{\pi_{t_0}, t_0>0\}$ with the property that if $\rho<1$ there exists a policy in the family under which the system is stable. Notice that the result is weaker than the result in the previous section where it is shown that for all $\rho<1$ the system is stabilized by CFP. The results in this section are obtained under the following statistical assumption.

S2. The interarrival and the service times are i.i.d., and they have all moments.

The above condition is satisfied by most service time distributions in practice and by all distributions with finite support. The class of policies $\{\pi_{t_0}, t_0>0\}$ is defined as follows.

Policy π_t

STEP 0. Initialize $i=0$.

STEP 1. At time $\tau_i=it_0$ a sequence of demands is determined by *CFA* for the configuration of the demands in $\mathcal{A}$ at that time.

STEP 2. The demands are served in the order that they have been selected in Step 1. If the service of all those demands finishes before time τ_{i+1} then the server serves the remaining demands in an arbitrary order. At time τ_{i+1}, if the server is busy the served demand is preempted, i is increased by 1, and the process is repeated from Step 1.

The following theorem states the stability properties of the class of policies π_t.

Theorem 3. *If $\rho<1$ and the assumption S2 holds, then there exists $t_0>0$ such that under policy π_{t_0} the system is stable in the sense*

$$\sup_t E[Q(t)]<\infty. \quad (17)$$

Note that the property (17) of the queue length process is stronger than (1). Therefore it implies stability.

The proof of the theorem will be given later after some preliminary results. We will analyze first the process Q = $\{Q_i\}$ where $Q_i = Q(\tau_i)$. We need some results on the drift analyses of random sequences with general statistics obtained by Hajek (1982). They are briefly described in the following for completeness.

4.1. Background on Drift Analysis of Random Sequences with General Statistics

Consider the sequence of random variables $\{Q_k\}_{k=1}^{\infty}$ on a probability space $(\Omega, \mathcal{F}, P)$ which are adapted on an increasing family $\{\mathcal{F}_k\}_{k=1}^{\infty}$ of subfields of $\mathcal{F}$. The drift of that sequence is by definition $E[Q_{k+1} - Q_k|\mathcal{F}_k]$. Two conditions on the drift are given next which imply a stability property of the random sequence.

C1. For some b and ϵ_0 with $-\infty \leq b < \infty$ and $\epsilon_0 > 0$,

$$E[Q_{k+1} - Q_k|Q_k > b; \mathcal{F}_k] \leq -\epsilon_0 \quad \text{for } k \geq 0 .$$

The second condition on the drift requires the following definitions. Let $\mathcal{G}$ be a sub-σ-field of $\mathcal{F}$; then we say that a random variable Z stochastically dominates a random variable Y given $\mathcal{G}$, written $(Y|\mathcal{G}) < Z$, if $P(Y > c|\mathcal{G}) \geq P(Z > c)$ for $-\infty < c < \infty$.

C2. $(|Q_{k+1} - Q_k||\mathcal{F}_k) < Z$ for all $k \geq 0$ and $Ee^{\lambda Z} = D < \infty$.

Choose constants η and ρ such that

$$0 < \eta \leq \lambda,\ \eta < \frac{\epsilon_0}{c},\ \rho = 1 - \epsilon_0\eta + c\eta^2,$$

where

$$c = \frac{Ee^{\lambda Z} - (1 + \lambda EZ)}{\lambda^2} .$$

The following is shown in Hajek.

Theorem 4. *When the random sequence $\{Q_k\}_{k=1}^{\infty}$ satisfies conditions C_1 and C_2 we have*

$$E[e^{\eta Q_k}|\mathcal{F}_0] \leq e^{\eta Q_1} + De^{\eta a}, \tag{18}$$

and $0 < \rho < 1$.

4.2. Proof of Theorem 3

Using the results of the previous section we show the following.

Lemma 3. *If $\rho < 1$ there exists a t_0 such that under policy π_{t_0} we have*

$$\sup_{i \geq 1} E\{Q_i\} \leq B, \tag{19}$$

for some B which may depend on ρ.

Lemma 3 follows easily from Theorem 4 if we verify the Conditions C1, C2 for the process Q. The next lemma shows Condition C1.

Lemma 4. *For all scheduling policies π_{t_0}, $t_0 > 0$ there exists some $\lambda > 0$ and random variable X such that*

$$(|Q_{i+1} - Q_i||\mathcal{G}_1^i) < X,$$

$$Ee^{\lambda X} = D < \infty, \tag{20}$$

where $\mathcal{G}_1^i$ denotes the σ-field generated from the random variables $Q_1, \ldots, Q_i$.

Proof. Let A_i, S_i be the number of arrivals and departures, respectively, at the time interval $[it_0, (i+1)t_0)$. Then we have

$$Q_{i+1} - Q_i = A_i - S_i ;$$

therefore

$$|Q_{i+1} - Q_i| \leq A_i + S_i \quad \text{a.s.},$$

and

$$(|Q_{i+1} - Q_i||\mathcal{G}_1^i) < A_i + S_i . \tag{21}$$

Let $A(t)$ be a version of the arrival process such that the first arrival occurs at time $t = 0$ and $S(t)$ a renewal process, independent of $A(t)$, with the time between two consecutive points distributed with the service time distribution. Clearly we have $A_i < A(t_0)$, $S_i < S(t_0)$. If we take $X = A(t_0) + S(t_0)$ we have

$$(|Q_{i+1} - Q_i||\mathcal{G}_1^i) < X .$$

We will show now that $Ee^{\lambda X} < \infty$ for some $\lambda > 0$. Since $A(t_0)$ and $S(t_0)$ are independent we have

$$Ee^{\lambda X} = Ee^{\lambda A(t_0)}Ee^{\lambda S(t_0)} . \tag{22}$$

In view of (22) in order to show (20) it is enough to show that for some λ'

$$Ee^{\lambda' A(t_0)} < \infty .$$

Then for some λ'', $Ee^{\lambda'' S(t_0)} < \infty$ as well, and (20) holds for $\lambda = \min\{\lambda', \lambda''\}$. We have

$$Ee^{\lambda A(t_0)} \leq \sum_{k=0}^{\infty} e^{\lambda k} P[A(t_0) \geq k] . \tag{23}$$

Let $N_j = \min\{t: A(t) \geq j\}$, $F_j(t) = P[N_j \leq t] = P[A(t) \geq k]$. It is shown in Karlin and Taylor (1975, p. 181–182) that for any integers m, r, n we have

$$F_{nr+m}(t) \leq [F_r(t)]^n F_m(t), \quad 0 \leq m \leq r - 1 . \tag{24}$$

From (23) and (24) we have

$$Ee^{\lambda A(t_0)} \leq \sum_{m=0}^{r-1} F_m(t_0)e^{\lambda m} \sum_{n=0}^{\infty} e^{\lambda rn}[F_r(t_0)]^n,$$

$$0 \leq m \leq r - 1 ,$$

hence it is enough to show

$$\sum_{n=0}^{\infty} (e^{\lambda r}F_r(t_0))^n < \infty . \tag{25}$$

Clearly for any nontrivial interarrival distribution, there exists an r_0 such that $F_{r_0}(t_0) < 1$. Chose $\lambda < -1/r \ln F_r(t_0)$. Then

$$e^{\lambda r_0} F_{r_0}(t_0) < 1 ,$$

and (25) follows. □

In the following lemma we verify Condition C2.

Lemma 5. *If $\rho < 1$ then there exists t_0, b, $\epsilon_0 > 0$ such that when the system is scheduled under π_{t_0} we have*

$$E[Q_{i+1} - Q_i | Q_i \geq b; \mathcal{G}_1^i] \leq -\epsilon_0 . \tag{26}$$

Proof. Let A_i, S_i be the number of arrivals and departures, respectively, in the time interval $[it_0, (i+1)t_0)$. Then we have

$$E[Q_{i+1} - Q_i | Q_i > b; \mathcal{G}_1^i] = E[A_i | Q_i > b; \mathcal{G}_1^i] - E[S_i | Q_i > b; \mathcal{G}_1^i] . \tag{27}$$

Let $A(t)$ be a version of the arrival process. Clearly we have

$$E[A_i | Q_i > b; \mathcal{G}_1^i] \leq E[A(t_0)] . \tag{28}$$

Let $S(t)$ be a renewal process, independent of $A(t)$, with the time interval between points j and $j+1$ been equal to the service time of the jth demand served after the time it_0. We can see, after some thought, that the following holds a.s.

$$S_i \geq \min\{\tfrac{1}{2}\log Q_i, S(t - 2\sqrt{2}a)\} .$$

Therefore,

$$E[S_i | Q_i > b; \mathcal{G}_1^i] \geq E[\min\{\tfrac{1}{2}\log Q_i, S(t - 2\sqrt{2}a)\} | Q_i > b; \mathcal{G}_1^i] . \tag{29}$$

By manipulating the right hand side of (29) we get

$$\begin{aligned} &E[\min\{\tfrac{1}{2}\log Q_i, S(t - 2\sqrt{2}a)\} | Q_i > b; \mathcal{G}_1^i] \\ &= E[S(t - 2\sqrt{2}a) | \tfrac{1}{2}\log Q_i \geq S(t - 2\sqrt{2}a); Q_i > b; \mathcal{G}_1^i] \\ &\quad \cdot P[\tfrac{1}{2}\log Q_i \geq S(t - 2\sqrt{2}a) | Q_i > b; \mathcal{G}_1^i] \\ &\quad + \tfrac{1}{2}\log Q_i P[S(t - 2\sqrt{2}a) \geq \tfrac{1}{2}\log Q_i | Q_i > b; \mathcal{G}_1^i] . \end{aligned} \tag{30}$$

From (29), (30) and using the analysis of $E[S(t - 2\sqrt{2}a)]$ as a sum of conditional expectations we get

$$\begin{aligned} E[S_i | Q_i > b; \mathcal{G}_1^i] &\geq E[S(t - 2\sqrt{2}a)] \\ &- E[S(t - 2\sqrt{2}a) | \tfrac{1}{2}\log Q_i < S(t - 2\sqrt{2}a); \\ &\quad Q_i > b; \mathcal{G}_1^i]P \\ &\cdot [\tfrac{1}{2}\log Q_i < S(t - 2\sqrt{2}a) | Q_i > b; \mathcal{G}_1^i] . \end{aligned} \tag{31}$$

By replacing in (27) from (31) we get

$$\begin{aligned} &E[Q_{i+1} - Q_i | Q_i > b; \mathcal{G}_1^i] \\ &\leq E[A(t_0)] - E[S(t_0 - 2\sqrt{2}a)] \\ &\quad + E[S(t - 2\sqrt{2}a) | \tfrac{1}{2}\log Q_i < S(t - 2\sqrt{2}a); \\ &\qquad Q_i > b; \mathcal{G}_1^i] \\ &\quad \cdot P[\tfrac{1}{2}\log Q_i < S(t - 2\sqrt{2}a) | Q_i > b; \mathcal{G}_1^i] . \end{aligned} \tag{32}$$

From the elementary renewal theorem we have

$$\lim_{t_0 \to \infty} \frac{EA(t_0) - ES(t_0 - 2\sqrt{2}a)}{t_0} = \lambda - \mu < 0 ,$$

therefore there exists some t_0 and $\epsilon > 0$ such that

$$EA(t_0) - ES(t_0 - 2\sqrt{2}a) = -\epsilon < 0 . \tag{33}$$

From (32) and (33) and the fact that the last term of the sum in (32) converges to 0 as b goes to infinity, we conclude that there exists b large enough, t_0 and $\epsilon_0 > 0$ such that (26) holds. □

From Lemma 5 the theorem follows easily.

Proof of Theorem 3. For an arbitrary time t' let τ_i, τ_{i+1} be the decision time instants such that $\tau_i \leq t' < \tau_{i+1}$. Using Lemma 4 we get

$$\begin{aligned} E[Q(t')] &= E[Q(\tau_i)] + E[A(t')] - E[S(t')] \\ &\leq E[Q(\tau_i)] + E[A(t')] + E[S(t')] \\ &\leq B + A(t_0) + S(t_0) + 2 \leq \infty , \end{aligned}$$

where $A(t)$ is a version of the arrival process and $S(t)$ is a renewal process, independent of $A(t)$, with the time interval between points j and $j+1$ been equal to the service time of the jth demand served after the time it_0. □

5. ADAPTIVE PARTITIONING POLICIES

A partitioning policy operates as follows. The service region is divided in l^2 subregions by a grid of l vertical and l horizontal parallel equidistant lines (Figure 3) where l is a parameter of the policy. The server visits the subregions one after the other in a prespecified manner, as it is depicted in Figure 3. The server moves from one subregion to the next after it serves all the demands in the subregion. The partitioning policy has been proposed by Bertsimas and Van Ryzin (1991a); its performance was analyzed, and it was shown that as long as $\rho < 1$ there is a partitioning policy with a specific parameter l for which the system is stable; the parameter is determined based on the arrival and service rates.

In this section we present a class of adaptive partitioning policies where the grid changes during the operation of the system. The parameter l, that determines the resolution of the grid is computed at the times instances at which the grid changes, based on the number of service demands on the plane at that time. These policies stabilize the system for every $\rho < 1$ without requiring knowledge of the arrival

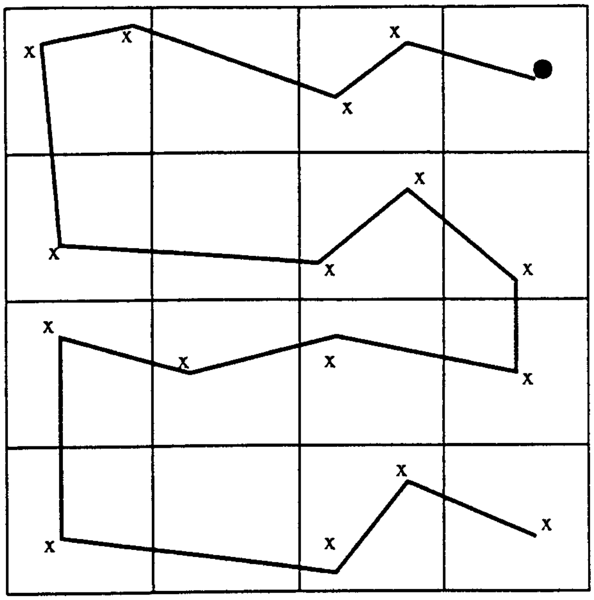

Figure 3. The tour followed by the server during one cycle of the adaptive partitioning policy, where the configuration of the demands on the plane at the beginning of the cycle is as depicted above.

and service rates λ and μ. The stability result holds under general location processes and Poisson arrivals.

The adaptive partitioning policies operate in cycles. The ith cycle lasts from time τ_{i-1} to τ_i. During each cycle the policy behaves like a partitioning policy. The grid that determines the subregions for the implementation of the partitioning policy changes from cycle to cycle. Let l_i be the parameter of the grid in the ith cycle; also let $Q_i = Q(\tau_i)$ and g is a function g: $Z^+ \rightarrow Z^+$. The policy operates as follows.

Cycle i.

A grid is considered on the region $\mathcal{A}$ with parameter $l_i = g(Q_i)$. During the ith cycle the server visits each subregion determined by the grid once and following the route depicted in Figure 3. In each subregion the server servers all the demands that were present in the subregion in the beginning of the ith cycle and not any demand that arrived after τ_{i-1}. After each subregion has been visited the next cycle starts.

The function g distinguishes the different adaptive partitioning policies. As long as g satisfies

$$\lim_{t\to\infty} \frac{g(t)}{t} = 0, \quad \lim_{t\to\infty} g(t) = \infty, \tag{34}$$

the system is stable as is stated in the following theorem. Let π^g denote the adaptive policy that uses function g to determine the grid density at the beginnings of the cycle.

Theorem 5. *When the system is operated under policy π^g, the function g satisfies Conditions* (34) *and the statistics satisfy assumption $A1$ then for every $\rho < 1$*

$$\limsup_{t\to\infty} \frac{1}{t}\int_0^t EQ(s)\,ds < \infty,$$

and the system is stable.

The proof of the Theorem 5 is similar to that of Theorem 2. We just prove here a drift condition on the queue length which is stated in Lemma 6. The rest of the proof is omitted.

Lemma 6. *When the system is operated under policy π_g, the function g satisfies Condition* (34), *$\rho < 1$ and the statistics satisfy assumption $A1$ then there exists B, $\epsilon > 0$ such that*

$$E[Q_{i+1}^2 - Q_i^2|Q_i] < -\epsilon Q_i^2 \quad \text{if } Q_i > B. \tag{35}$$

Proof. We have

$$E[Q_{i+1}^2|Q_i] = E[(A(\tau_{i+1}) - A(\tau_i))^2|Q_i] = \lambda^2 E[(\tau_{i+1} - \tau_i)^2|Q_i]. \tag{36}$$

Notice that $\tau_{i+1} - \tau_i$ is equal to the total time spent in the service of the Q_i demands plus the time spent in traveling from demand to demand. The traveling time in a cycle consists of two parts. The first is the time spent in traveling from subregion to subregion, and the second is the time spent traveling from demand to demand within each subregion. The first time is less than or equal to $2\sqrt{2}al_i$. For the second time note that, irrespective of the configuration of the points in the region, the time for going from one point to another in a straight line within a subregion is bounded by the length of the diagonal of the subregion which is $\sqrt{2}a/l_i$. Therefore the total time spent in traveling within the subregions is bounded above by $Q_i\sqrt{2}a/l_i$ and the total traveling time is bounded as

$$\textit{total traveling time in the } i+1 \textit{ cycle} \le 2\sqrt{2}al_i + Q_i\frac{\sqrt{2}a}{l_i}. \tag{37}$$

Let S_j be the service time of the jth customer served in the time interval $\tau_{i+1} - \tau_i$. We have

$$E[(\tau_{i+1} - \tau_i)^2|Q_i] \le E\left[\left(\sum_{j=1}^{Q_i} S_j + 2\sqrt{2}al_i + Q_i\frac{\sqrt{2}a}{l_i}\right)^2\Big|Q_i\right] \le Q_i^2 E\left[\left(\frac{\sum_{j=1}^{Q_i} S_j}{Q_i} + 2\sqrt{2}a\frac{g(Q_i)}{Q_i} + \frac{\sqrt{2}a}{g(Q_i)}\right)^2\Big|Q_i\right]. \tag{38}$$

Given that $E[S_j^2] < \infty$, as Q_i increases, the term in the parenthesis in (37) converges to μ^{-1} almost surely from the law of large numbers and because $g(\,)$ satisfies conditions (34). Hence the expectation in (38) converges to μ^{-2}, and clearly there exists $\rho_0 < 1$ such that for B large enough we get

$$\lambda^2 E[(\tau_{i+1} - \tau_i)^2|Q_i] \le \rho_0^2 Q_i^2 \quad \text{if } Q_i > B. \tag{39}$$

From (36) and (39) we conclude that (35) holds for $\epsilon = \rho_0^2 - 1$. □

6. CONCLUSIONS AND FURTHER RESEARCH

In this paper we considered the problem of routing a server that provides service to demands arising in a region of the plane. We presented an algorithm, that, given N arbitrarily located points on the plane, computes a tour that passes through $\log N/2$ of these points and has total length upper bounded by a constant. Based on that algorithm we obtained the policy CFP, which stabilizes the system, as long as $\rho < 1$, for any location process. CFP is independent of the system statistics. Then we studied the system with renewal arrival process. We showed the existence of a parametrized policy which for appropriate choices of the parameter stabilizes the system as long as $\rho < 1$. That policy, unlike CFP, depends on the system statistics. No statistical assumptions were made for the locations of the demands on the plane and the stability result holds for every location pattern of the demands. An adaptive version of the partitioning policies was given. The adaptive policy becomes independent of the statistics of the system without deterioration of the throughput.

This work leaves several problems open for further investigation. The analysis of the delay induced by the CFP is one of them. The property of the CFA algorithm to route the server to the region of the plane with the higher concentration of demands has as a result the nice stability properties of the CFP. Due to the complicated nature of the policy, though, a delay analysis based on standard queueing techniques seems to be infeasible, and a different approach is needed. In addition the delay will depend on the location distribution and its analysis needs further, extensive investigation. The study of variations of CFP is an interesting problem, too. For example, at the end of a cycle of CFP, instead of reapplying the CFA in the whole plane, the server may backtrack to the square most recently visited and start serving requests there. A classification of these alternatives, with respect to the delay that they achieve, is of interest. Finally we saw that there is a large class of adaptive versions of the partitioning policies that achieve maximum throughput. These policies differ on the way they adjust the grid that partitions the plane. Some further investigation, to distinguish adaptive policies with superior performance as far as the delay is concerned, is in order.

REFERENCES

ASMUSSEN, S. 1987. *Applied Probability and Queues*. J. Wiley and Sons, New York.

BATTA, R., R. C. LARSON, AND A. R. ODONI. 1988. A Single Server Priority Queueing Location Model. *Networks*, **18**, 87–103.

BERTSIMAS, D. J., AND G. VAN RYZIN. 1991. A Stochastic and Dynamic Vehicle Routing Problem in the Euclidean Plane. *Opns. Res.* **39**, 601–615.

BERTSIMAS, D. J., AND G. VAN RYZIN. 1993a. Stochastic and Dynamic Vehicle Routing in the Euclidean Plane with Multiple Capacitated Vehicles. *Opns. Res.* **41**, 60–76.

BERTSIMAS, D. J., AND G. VAN RYZIN. 1993b. Stochastic and Dynamic Vehicle Routing in with General Demand and Interarrival Time Distributions. *Adv. Appl. Prob.* **25**, 947–978.

JAILLET, P. 1988. A Priori Solution of a Traveling Salesman Problem in Which a Random Subset of the Customers Are Visited. *Opns. Res.* **36**, 929–936.

KARLIN, S., AND H. M. TAYLOR. 1975. *A First Course in Stochastic Processes*. Academic Press, New York.

KARP, R. 1977. Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman in the plane. *Math. O. R.* **2**, 209–224.

PSARAFTIS, H. N. 1988. Dynamic Vehicle Routing Problems. In *Vehicle Routing: Methods and Studies*. B. Golden and S. Assad (eds.). North-Holland, Amsterdam.

THOMPSON, P. M., AND H. N. PSARAFTIS. 1993. Cyclic Transfer Algorithms for Multivehicle Routing and Scheduling Problems. *Opns. Res.* **41**, 935–946.